

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 18/Jan/2002		2. REPORT TYPE MAJOR REPORT		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE SCALABLE ELECTROMAGNETIC SCATTERING COMPUTATIONS				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) CAPT HASTRITER MICHAEL L				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) UNIVERSITY OF ILLINOIS AT URBANA				8. PERFORMING ORGANIZATION REPORT NUMBER CI02-9	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) THE DEPARTMENT OF THE AIR FORCE AFIT/CIA, BLDG 125 2950 P STREET WPAFB OH 45433				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Unlimited distribution In Accordance With AFI 35-205/AFIT Sup 1				DISTRIBUTION STATEMENT A: Approved for Public Release - Distribution Unlimited	
13. SUPPLEMENTARY NOTES					
<div style="font-size: 2em; font-weight: bold; margin-bottom: 10px;">20020204 094</div>					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 4	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code)

Scalable Electromagnetic Scattering Computations

Sanjay Velamparambil*, Weng Cho Chew, and M. Larkin Hastriter

Center for Computational Electromagnetics
Electromagnetics Laboratory
Department of Electrical and Computer Engineering
University of Illinois, 1406 W. Green St., Urbana, IL 61801-2991, USA
{velampar, w-chew, hastrite}@uiuc.edu

1 Introduction

The success of MLFMA in solving large scale problems has naturally led to efforts in parallelizing the algorithm. Several academic and industrial research groups have made significant progress in their attempts [1, 2, See references therein]. We have shown that the parallel implementation of MLFMA developed at the University of Illinois, called *ScaleME*, has excellent scaling properties [2].

In this paper, we summarize the results of our efforts in developing a scalable distributed memory fast electromagnetic integral equation solver. The massively parallel scattering code, called LSSP, is based on a Galerkin discretization of the combined field integral equation (CFIE) using the RWG basis functions. The resulting matrix equation is then solved using a parallel GMRES solver [3]. It uses the distributed memory parallel MLFMA library called *ScaleME* for evaluating matrix-vector products and Message Passing Interface (MPI) for inter-processor communication. Furthermore, once the solution is obtained, the bistatic RCS is computed using a parallelized, MLFMA based, far field evaluation algorithm [1]. The objective of this paper is to present some recent results demonstrating the scalability of the code for solving realistic problems.

2 Summary of Parallelization of MLFMA

The basic idea in parallelizing MLFMA can be described as follows: *for the top (coarse) few levels, replicate the boxes in every processor, but split the far field patterns equally among all processors. For the finer levels, divide the boxes at each level equally among the processors.*

The levels that are replicated in every processor are called “shared” levels, and the levels for which the grain size is retained to be a box are called “distributed” levels. Clearly, this scheme divides the tree horizontally into different layers, each layer consisting of one or more levels. These overlapping layers are the *shared layer*, the *transition layer*, and the *distributed layer*. Each layer has distinct communication and computational behaviors. For instance, for the distributed layer, communications are necessary during all three phases, viz. the aggregation, translation, and disaggregation phases. However, for the shared and transition layers, no communication is necessary during the translation phase. Furthermore, during the aggregation and disaggregation phases, these two layers require communication of partial radiation/receiving patterns. In fact, during the aggregation and the disaggregation phases, the parallel interpolation and antepolation operations are required.

However, since each processor has only N_s/p samples of the far field pattern, the maximum length of the messages is bounded by the same amount. Indeed, using a local interpolation/antepolation scheme, the length of the message can be reduced from $O(N)$ to $O(\sqrt{N})$ [2].

3 Parallel MLFMA based RCS computations

The evaluation of RCS after solving the matrix equation is computationally expensive. To see this, note that evaluating the RCS in a single direction requires $O(N)$ operations. For large scale problems, the number of directions in which the RCS is sought on one plane cut is $O(\sqrt{N})$, thereby requiring a total evaluation time of $O(N^{1.5})$. Furthermore, the proportionality constant is rather large owing to the various geometric computations involved.

In this section, we briefly discuss a method by which the bistatic RCS can be evaluated rapidly using parallel computers. The bistatic RCS can be written as,

$$\sigma(\theta, \phi) = \frac{\eta^2}{4\pi} \left| \sum_{i=1}^N \alpha_i k(\bar{\mathbf{I}} - \hat{s}\hat{s}) \cdot \int_S \mathbf{f}_i(\mathbf{r}') e^{-ik\hat{s} \cdot \mathbf{r}'} dS' \right|^2 \quad (1)$$

where \mathbf{f}_i is the i -th basis function and α_i the corresponding coefficient. Also, we have set $\hat{s} = (1, \theta, \phi)$. For a given point $\mathbf{c} \in \mathbb{R}^3$,

$$\mathbf{F}_m(\hat{s}) = k(\bar{\mathbf{I}} - \hat{s}\hat{s}) \cdot \int_S \mathbf{f}_m(\mathbf{r}') e^{ik\hat{s} \cdot (\mathbf{c} - \mathbf{r}')} dS', \quad (2)$$

is the radiation pattern of the m -th basis function. By setting $\mathbf{c} = 0$, we see that the summand in Equation (1) is the *radiation pattern* of the m -th basis function with respect to the origin.

Since we have the radiation patterns of the basis functions with respect to the finest level of MLFMA tree, we can compute the sum in Equation (1) using the upward pass of MLFMA. However, the upward pass needs to be supplemented with an interpolation at the top-most level. This interpolation must be parallelized efficiently.

Depending on whether the algorithm uses shared levels or not, the parallelization requires two different approaches. In the absence of shared levels, parallelization involves only a global reduction operation after the interpolation to the root box. When shared levels are present, the algorithm has to take into account sparsity in the data structures and thus is more involved. The details of these methods are described in [1].

4 Numerical Results

We have verified the correctness and the accuracy of the code by comparing with analytical solutions for perfectly conducting spheres, as well as by comparing with other results available from literature. Here we present two sets of results demonstrating the scalability of the methods employed. Let p be the number of processors and T_1 and T_p be the time taken by the algorithm on one processor and on p processors, respectively. Then, the parallel efficiency is $\eta = T_1/(pT_p)$.

First, we demonstrate the parallel efficiency of MLFMA for evaluating matrix vector products. For this, we consider the scattering from a 15λ perfectly conducting cube modeled using 294,912 unknowns. The parallel efficiency is plotted in Figure 1. The figure demonstrates that the shared levels improves the efficiency significantly as the number of processors is increased. This behavior is consistent with the theoretical analysis [2]. However, it may also be noted that by increasing the number of shared levels from 2 to 3 does not improve the performance very much. In fact, we have observed that depending on the geometry, there is a range of shared levels for which the performance is more or less the same. Finally, the “superlinear” performance exhibited for a small number of processors is a result of machine dependent memory hierarchy.




Figure 1: Parallel efficiency for the case **CUBE-15λ**. SLEV refers to the finest shared level used in the simulation.

Next, we demonstrate the scalability of the code with respect to the number of unknowns using a full-size, fictitious aircraft, referred to as VFY-218. We study the model with 625,626 unknowns at 2 GHz and 2,464,536 unknowns at 4 GHz. The total run time for the first case was about 2 hours and 13 minutes and about 5 hours for the second. The time for evaluating the matrix-vector product and for evaluating the RCS for 1800 directions are given in Table 1. The results demonstrate very good scaling.

Num. Proc	MatVec Time(s)		RCS Eval. Time(s)	
	2 GHz	4 GHz	2 GHz	4 GHz
16	59.55		10.71	
32	30.87	111.16	6.40	20.02
64		64.41		11.35

Table 1: Demonstration of scalability with respect to problem size for the full scale model of an aircraft, VFY-218. At 2 GHz, the number of unknowns $N = 625,626$ and at 4 GHz, $N = 2,464,536$.

Finally, we present the bistatic RCS of the aircraft model at 8 GHz. The four-de-force simulation involved 10,186,446 unknowns and we used a 10-level MLFMA, and 126 processors of an SGI Origin 2000 supercomputer. The total solution time was 7 hours and 25 minutes and each matrix-vector product evaluation required 119 seconds. The latter results again shows the scaling with respect to the problem size. The bistatic RCS for the vertical polarization with an incident direction of (90,90) is shown in Figure 2.




Figure 2: Bistatic RCS of VFY-218 at 8 GHz. $N = 10,186,446$.

5 Conclusions

The objective of the paper was to present a brief summary of the scalable parallel code we have developed for electromagnetic scattering computations. We have presented representative results demonstrating the excellent scalability obtained. More detailed results will be presented at the conference.

References

- [1] S. Velampambil, W. C. Chew, and J. Song, "10 million unknowns: Is it *that* big?," Tech. Rep. 15-01, Center for Computational Electromagnetics, 2001. Submitted to IEEE Antennas and Propagation Magazine.
- [2] S. Velampambil and W. C. Chew, "Parallelization of multilevel fast multipole algorithm on distributed memory computers," Tech. Rep. 13-01, Center for Computational Electromagnetics, University of Illinois at Urbana-Champaign, 2001. Under the review of Parallel Computing.
- [3] V. Frayssé, L. Giraud, and S. Gratton, "A set of GMRES routines for real and complex arithmetics," Tech. Rep. TR/PA/97/49, CERFACS, France, 1997. <http://www.cerfacs.fr/algor/Softs/index.html>.